



Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

A sequential perspective on searching for static targets

Kurt E. Wilson, Roberto Szechtman*, Michael P. Atkinson

Operations Research Department, Naval Postgraduate School, Monterey, CA 93943, United States

ARTICLE INFO

Article history:

Received 4 January 2010

Accepted 25 May 2011

Available online xxxx

Keywords:

Applied probability

OR in military

Sequential analysis

ABSTRACT

We present a sequential approach to detect static targets with imperfect sensors, which range from tower-mounted cameras to satellites. The scenario is operationally relevant to many military, homeland security, search and rescue, environmental engineering, counter-narcotics, and law enforcement applications. The idea is to stop the search as soon as there is enough probabilistic evidence about the targets' locations, given an operator-prescribed error tolerance, knowledge of the sensors' parameters, and a sequence of detection signals from the sensors. By stopping the search as soon as possible, we promote efficiency by freeing up sensors and operators to perform other tasks. The model we develop has the added benefits of decreasing operator workload and providing negative information as a search progresses.

Published by Elsevier B.V.

1. Introduction

Today's operational planners and sensor operators face numerous challenges inherent to the complex environments that shape their space of operation. These challenges are further magnified by scarce resources, imperfect information, and operator task overload.

The time critical nature of the command decisions that serve as milestones throughout the Find, Fix, Track, Target, and Engage (F^2T^2E) process further exacerbate the situation. Defense planners must strive to develop and incorporate new, efficient procedures to allocate scarce resources in many different complex environments. Any efficiency gained within the F^2T^2E chain, however small, may have a compound effect over time on overall operational readiness because this will free up assets to perform other time-sensitive, critical sensing actions, as well as decrease operator workload and capitalize upon negative information. Such negative information could be utilized to find where targets are *not* located, and may help determine areas to set up certain operations or paths through the environment that are free of hostile forces.

In this article we consider a scenario with multiple fixed-sensors and multiple static targets in discrete-time and discrete-space. The sensors may range from tower-mounted cameras, to Unmanned Aircraft Systems (UASs), to satellites, and the targets under consideration do not react to any sensing action. The scenario is operationally relevant to many military, homeland

security, search and rescue (SAR), environmental engineering, counter-narcotics, and law enforcement applications. UASs have been used in Iraq and Afghanistan to search for Improvised Explosive Devices (IEDs), insurgent safe houses, suspected weapons caches, and mortar points of origin [11,12]. Other relevant applications include searching for downed aircraft or life rafts, detecting illegal drug harvesting and processing operations, patrolling border infiltration points, and tracking flora and fauna counts in biological environments.

We formulate a model to locate static targets of interest (TOIs) hidden within an area of interest (AOI). As in reality, our model allows the analyst to contend with the fact that the search sensors are imperfect; i.e., the sensors may declare fewer or more targets than are actually present on a particular search attempt. The idea is to stop the search as soon as there is enough probabilistic evidence about the TOIs' locations, given an operator-prescribed error tolerance, knowledge of the sensors' parameters, and a sequence of detection signals from the sensors.

The AOI for the scenario is comprised of a grid of discrete, non-overlapping area-cells (ACs). The area-cells might be defined by geo-political borders, terrain features, or some arbitrary grid system of tactical significance to the operator, and need not be uniform in size nor shape. Each cell is characterized by the number of sensors (known), the sensors' operational parameters (known), and the number of targets (unknown).

The sensor parameters are the conditional probabilities of returning each possible detection signal given each possible number of actual TOIs in that area-cell. More specifically, when the operator makes an investigation into an area-cell, the sensor returns a detection signal corresponding to the number of TOIs seen by the sensor with some probability that depends on the actual

* Corresponding author. Tel.: +1 831 656 3311.

E-mail addresses: kurt.e.wilson@afghan.swa.army.mil (K.E. Wilson), rszechtm@nps.edu (R. Szechtman), mpatkins@nps.edu (M.P. Atkinson).

Report Documentation Page			Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.				
1. REPORT DATE 2011		2. REPORT TYPE		3. DATES COVERED 00-00-2011 to 00-00-2011
4. TITLE AND SUBTITLE A Sequential Perspective On Searching For Static Targets			5a. CONTRACT NUMBER	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)			5d. PROJECT NUMBER	
			5e. TASK NUMBER	
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School, Operations Research Department, Monterey, CA, 93943			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited				
13. SUPPLEMENTARY NOTES European Journal of Operational Research, preprint, 2011				
14. ABSTRACT We present a sequential approach to detect static targets with imperfect sensors, which range from tower-mounted cameras to satellites. The scenario is operationally relevant to many military, homeland security, search and rescue, environmental engineering, counter-narcotics, and law enforcement applications. The idea is to stop the search as soon as there is enough probabilistic evidence about the targets?locations, given an operator-prescribed error tolerance, knowledge of the sensors? parameters, and a sequence of detection signals from the sensors. By stopping the search as soon as possible, we promote efficiency by freeing up sensors and operators to perform other tasks. The model we develop has the added benefits of decreasing operator workload and providing negative information as a search progresses.				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 10
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified		

(i.e., the *ground truth*, which is unknown) number of TOIs in that area-cell.

To efficiently determine the targets' locations subject to the operator-prescribed error tolerance, we develop a sequential eliminating procedure [14]. A sequential eliminating procedure attempts to isolate, from among several candidate configurations, one particular desired configuration—the objective. During a particular stage of a sequential eliminating procedure, all candidate configurations are examined and ranked in order of their likelihood of producing the sequence of observed signals up to that stage.

Any configuration whose likelihood, when compared with the configuration of maximum likelihood, exceeds a particular threshold (which depends on the user's error tolerance) is permanently eliminated from the set of candidates. If no configuration exceeds the threshold during a particular stage, then all those configurations remain in the set of candidates. The procedure advances to the next stage, using the updated candidate set. The process continues until only one configuration remains in the candidate set, and that configuration is declared the winner. In our case, the configurations are the ways the TOIs can be located in the area-cells of the AOI, and the winner is the *determined configuration*. We designate the actual location of the targets in the area-cells to be the *ground truth configuration* (GTC).

Search theory [1] traces its roots to the pioneering work of Koopman [5]. For the search scenario we focus on, the objective is to locate targets within a finite number of cells [1]. In this case, searcher success is achieved by either detecting the targets, or, if the targets are not detected, by correctly guessing the cells containing the targets. Tognetti [17] and Kadane [4] treat the scenario of whereabouts search against a stationary target. Washburn [19] is a classical reference in search theory.

Siegmund's 1985 book [14] is the classic reference in sequential analysis, and deals primarily with sequential hypothesis testing and related problems of estimation. In many of these cases, a fixed-sample solution exists, but one can employ sequential methods to achieve greater efficiency in the solution. Siegmund presents a sequential test with the same power as a fixed-sample test and requires fewer observations [14]. Therefore, the sequential test has a reasonable claim to be regarded as more efficient [14].

While the work to date in selection using sequential eliminating procedures [9,10,20] has focused on isolating the *best* system – usually the one with a maximum unknown parameter value – our goal is instead to isolate one determined configuration. The desire is for the determined configuration to be the ground truth configuration. That is it correctly specifies the number of TOIs in each area-cell. We show that our sequential model provides determined configurations efficiently, while guaranteeing to meet the user-prescribed error tolerance.

Compared to existing sensor employment models (e.g. [6,13,7,15]), our approach does not consider moving targets, does not dynamically allocate the sensors (i.e., no decision is taken as to where the sensors look in each stage), and does not find optimal search paths (with or without restrictions on searching area-cells within a vicinity of the last searched area-cell). Delving into the last point, most recent models (e.g. [6,13,7]) employ optimization techniques (deterministic or stochastic) that yield search paths that are optimal in a certain sense (often, but not always, maximizing the expected number of detected targets) for a prescribed number of time periods or search effort. However, these models scale poorly and become intractable even for a relatively small number of TOIs, area-cells, and time periods under consideration. This occurs because the computational cost grows at least exponentially in the number of variables ([18]), which generally is $\#TOIs \times \#ACs \times \#time\ periods$. Some heuristics (e.g. [16]) have been proposed to overcome this difficulty, but their performance cannot be theoretically guaranteed. While our approach does provide

some benefits over existing methods, it too has limitations with regard to the size of problems it can be applied to. The algorithm can only be used in situations where there are a handful of TOIs (e.g. less than 5) because it must consider all possible configurations of TOIs in the area cells. There are many applications where this is the case (e.g. SAR scenarios, searching for an insurgent safe house) and our algorithm would provide an appropriate and effective approach.

Once again, this article presents a sequential perspective on imperfect sensor employment applicable to static targets that is easy to implement, is computationally tractable when there are a small number of TOIs for a larger class of problems than the optimization approaches currently being employed, and stops when the user prescribed probabilistic guarantees are met (and thus the number of time periods or search effort is an output of the model).

The article is organized as follows. Section 2 introduces the notation and key definitions. In Section 3 we present the sequential approach. Section 4 shows numerical illustrations of the model, and Section 5 closes the paper with the main conclusions.

2. Notation

In this section we introduce the notation that will be employed throughout this article.

A : Number of area-cells.

M : Number of targets of interest.

m_i : Number of targets of interest in area-cell i . This value is unknown, and is what the analyst wishes to determine for $i = 1, \dots, A$.

m : The true configuration, $m = (m_1, m_2, \dots, m_A)$.

C : The set of feasible target configurations, formed by the elements $t = (t_1, \dots, t_A)$ non-negative and integer such that $\sum_{i=1}^A t_i = M$.

K : The number of feasible target configurations $K = |C| = \binom{M+A-1}{M}$.

$p_i(d|t_i)$: Conditional probability that the sensor in area-cell i returns a signal “ d targets” given that t_i TOIs are present there.

S_i : The sensor present in AC_i is completely characterized by the $(M+1) \times (M+1)$ matrix S_i . The value in the t th row and d th column of S_i is the probability $p_i(d|t)$. The matrix S_i is stochastic, so the elements of each row constitute a probability mass function.

$X_{i,1}, X_{i,2}, \dots$: Sequence of signals returned by the sensor in AC_i , independently and identically distributed (IID) with probability mass function $p_i(\cdot|m_i)$.

$\ell(x_{i,1}, \dots, x_{i,n}; t_i)$: For n IID signals from the sensor in AC_i , $x_{i,1} = x_{i,1}, x_{i,2} = x_{i,2}, \dots, x_{i,n} = x_{i,n}$, the likelihood of having t_i targets in area-cell i is

$$\ell(x_{i,1}, \dots, x_{i,n}; t_i) = \prod_{j=1}^n p_i(x_{i,j}|t_i).$$

3. Model

We describe the model in more detail in section 3.1. Our objective is to analyze the scenario of multiple targets in an area of interest with many cells. This produces many potential configurations. This case is difficult to analyze because knowledge about the presence/absence of targets in an area-cell yields insight about the presence/absence of targets in other area-cells; i.e., the number of targets in each area-cell is not independent. To gain insight into the problem, in Section 3.2 we examine the scenario where there are

only two possible configurations. In Section 3.3 we present the sequential eliminating algorithm, and in Section 3.4 we prove that this algorithm will finish in a finite number of steps with probability 1.

3.1. Background

We analyze a situation consisting of M TOIs located in A area-cells, and area-cell i contains m_i targets. The number of targets in each cell can be any value between 0 and M as long as $\sum_{i=1}^A m_i = M$; other constraint scenarios, such as $\sum_{i=1}^A m_i \leq M$, can be treated by enlarging the number of possible configurations. We denote the set of possible configurations of the M targets in the A area-cells as \mathcal{C} and therefore $|\mathcal{C}| = \binom{M+A-1}{M}$ (see page 38 of Feller [3] for a derivation). There is one sensor in each cell that provides an imperfect signal to the sensor operator about the number of targets in each cell. This model can allow for multiple sensors in each cell, but for notational simplicity we only analyze one sensor in each cell. The probability the sensor in area-cell i produces a signal that the area-cell contains d targets is $p_i(d|m_i)$. We construct a sensor matrix S_i in which the (t, d) element is given by $p_i(d|t)$. Each row of S_i , $p_i(\cdot|t)$, is a probability mass function (pmf). Thus the matrix S_i fully characterizes the sensor in AC_i . We assume that for all area-cells $1 \leq i \leq A$ the values $p_i(d|t)$ are known and that any two rows of the sensor matrix S_i have at least one different element.

Since the sensors are imperfect, it is desirable to frame the estimate with some level of confidence. To accomplish this, the operator prescribes an error tolerance. For example, an error tolerance of 5 percent would indicate that the operator is willing to accept that the model returns an acceptable configuration at least 95 percent of the time. In this paper, the only acceptable configuration is the ground truth configuration.

The objective for the sensor operator (and the one considered in this article) is to find the ground truth configuration subject to bounds on the probability of returning the wrong configuration. Naturally, the larger the error tolerance the operator is willing to accept, the more quickly the operator can expect the procedure to complete. Conversely, a small tolerance for error could mean many more search attempts are expended before the procedure terminates.

A key component of the sequential eliminating procedure is calculating the likelihood ratios between various configurations. The sequence of signals returned by the sensor in AC_i is an IID sequence of random variables $X_{i,1}, X_{i,2}, \dots$ drawn from the probability mass function $p_i(\cdot|m_i)$. If the realization of these signals is $X_{i,1} = x_{i,1}, X_{i,2} = x_{i,2}, \dots$, for ACs $i = 1, \dots, A$, then the joint likelihood function after n looks in each area-cell is given by

$$\ell(x_{ij}, 1 \leq i \leq A, 1 \leq j \leq n; t_1, \dots, t_A) = \prod_{i=1}^A \prod_{j=1}^n p_i(x_{ij}|t_i), \quad (1)$$

where $\{t_1, \dots, t_A\}$ corresponds to an arbitrary configuration. When the rows of each sensor matrix S_i are different, the likelihood function is maximized, as $n \rightarrow \infty$, by $t_1 = m_1, t_2 = m_2, \dots, t_A = m_A$; see Lehmann [8] for the general theory. For two arbitrary configurations $t^{(0)}$ and $t^{(1)}$, their likelihood ratio is

$$\frac{\ell(x_{ij}, 1 \leq i \leq A, 1 \leq j \leq n; t_1^{(0)}, \dots, t_A^{(0)})}{\ell(x_{ij}, 1 \leq i \leq A, 1 \leq j \leq n; t_1^{(1)}, \dots, t_A^{(1)})},$$

which is always a number in $(0, \infty)$ for finite n if $p_i(x_{ij}|t_i) > 0$.

3.2. Comparing two configurations

In this section we analyze the special case where there are two possible configurations, labeled Configuration 0 and Configuration 1, only one of which is the GTC. Configuration 0 has $t^{(0)} = \{t_1^{(0)}, t_2^{(0)}, \dots, t_A^{(0)}\}$ TOIs located among the area-cells, and Configuration 1 $t^{(1)} = \{t_1^{(1)}, t_2^{(1)}, \dots, t_A^{(1)}\}$ in the area-cells. If the ground truth configuration is Configuration 0, then the signals produced by the sensor in area-cell i will be distributed according to the pmf $p_i(\cdot|t_i^{(0)})$. If the ground truth configuration is Configuration 1, then the signals will be distributed according to the pmf $p_i(\cdot|t_i^{(1)})$. To guarantee that the algorithm terminates in finite time, we assume that at least one of the sensors' matrices has its $t_i^{(0)}$ and $t_i^{(1)}$ rows different; otherwise, the sequence of signals produced by the sensors would be probabilistically identical under both configurations.

We assume that all area-cells receive the same number of looks from their corresponding sensor, and the goal is to terminate the inspection when there is enough evidence that the error bounds are met, i.e., we are confident to within our error tolerances of saying that the TOIs are located in the area-cells according to one of the two configurations. The two possible errors are: (i) α is the probability that Configuration 1 is the determined configuration when the GTC is Configuration 0, and (ii) β is the probability that Configuration 1 is the determined configuration when the GTC is Configuration 0. This leads to the hypotheses

\mathcal{H}_0 : Configuration 0 is the GTC, and \mathcal{H}_1

: Configuration 1 is the GTC.

The likelihood that the signals $X_{ij} = x_{ij}$ are produced by Configuration 0 is (cf., Eq. (1))

$$\ell_n^{(0)} = \ell(x_{ij}, 1 \leq i \leq A, 1 \leq j \leq n; t_1^{(0)}, \dots, t_A^{(0)})$$

and the likelihood the signals are produced by Configuration 1 is

$$\ell_n^{(1)} = \ell(x_{ij}, 1 \leq i \leq A, 1 \leq j \leq n; t_1^{(1)}, \dots, t_A^{(1)}).$$

The likelihood ratio

$$\ell_n = \frac{\ell_n^{(1)}}{\ell_n^{(0)}} \rightarrow 0$$

with probability 1 as $n \rightarrow \infty$ if the GTC is Configuration 0 and $\ell_n \rightarrow \infty$ if the GTC is Configuration 1; see Lehmann [8]. This suggests that a judicious policy is to stop sampling when the likelihood ratio crosses an upper threshold and declare Configuration 1 the determined configuration or when the likelihood ratio crosses a lower threshold and declare Configuration 0 the determined configuration. This approach may lead to an incorrect determination, but its probability can be prescribed *ab initio* by the end-user.

Define the stopping time

$$N = \inf \{n \geq 1 : \ell_n \notin (A, B)\}$$

for threshold constants A, B such that $0 < A < B < \infty$. Then we

Reject \mathcal{H}_0 if $\ell_N \geq B$, and Accept \mathcal{H}_0 if $\ell_N \leq A$.

For $0 \leq \alpha, \beta \leq 1$ prescribed by the end-user, the error probabilities are

$$\begin{aligned} \text{Type I error : } P(\ell_N \geq B | \mathcal{H}_0) &= \alpha \text{ and Type II error : } P(\ell_N \\ &\leq A | \mathcal{H}_1) = \beta. \end{aligned}$$

Taking logarithms, we can see that we reject \mathcal{H}_0 if

$$\sum_{j=1}^N \sum_{i=1}^A \left(\log p_i(x_{ij}|t_i^{(1)}) - \log p_i(x_{ij}|t_i^{(0)}) \right) \geq \log B \quad (2)$$

and accept \mathcal{H}_0 if

$$\sum_{j=1}^N \sum_{i=1}^A \left(\log p_i(x_{ij}|t_i^{(1)}) - \log p_i(x_{ij}|t_i^{(0)}) \right) \leq \log A. \quad (3)$$

Siegmund [14] shows

$$\alpha = P(\ell_N \geq B | \mathcal{H}_0) \leq B^{-1}(1 - \beta) \quad (4)$$

and

$$\beta = P(\ell_N \leq A | \mathcal{H}_1) \leq A(1 - \alpha). \quad (5)$$

Hence, given operator defined tolerances α and β , by setting

$$B = \frac{1 - \beta}{\alpha}, \quad \text{and} \quad A = \frac{\beta}{1 - \alpha}, \quad (6)$$

we are guaranteed to satisfy the error probability constraints. Moreover, it is shown in Siegmund that if Eqs. (4) and (5) hold with equality, then this approach minimizes the expected number of looks until crossing either boundary [14]. Although (4) and (5) generally do not hold with equality, the algorithm is guaranteed to meet the error criteria.

Fig. 1 displays two possible sample paths for this sequential eliminating procedure. The two parallel dashed lines represent the bounds $\log(A)$ and $\log(B)$. A path between the bounds is still undetermined; thus the procedure continues until the path exits via one of the bounds (hence, this area is known as the *continuation region*). A path exiting the bound corresponding to $\log(B)$ results in a declaration that Configuration 1 is the determined configuration, whereas an exit via the $\log(A)$ bound results in a declaration that Configuration 0 is the determined configuration.

3.3. Sequential eliminating algorithm

We now consider the general scenario of an AOI consisting of $A > 1$ area-cells, and M targets located within the AOI. There are $K = \binom{M+A-1}{M}$ ways for the M targets to be placed in the A cells, leading to configurations $t^{(1)}, t^{(2)}, \dots, t^{(K)}$. Assume, without loss of generality, that the GTC is Configuration 1.

An incorrect determination occurs when the determined configuration is not Configuration 1. We call this event an *ICD* (for incor-

rect determination), and the desire is for $P(ICD) \leq \alpha$, for some $\alpha \in (0, 1)$ pre-specified by the operator.

Let C_0 be the initial set of candidate configurations; initially all K configurations are candidates to contain the target, so that $C_0 = C$. Let ℓ_n^* be the largest likelihood at stage n . We drop a configuration from consideration when there is sufficient evidence that it is not the correct configuration, i.e., when we are confident to within our error tolerance of saying that the targets are not located according to that particular configuration. This suggests eliminating configuration r when $\ell_n^* / \ell_n^{(r)} \geq B$, where $B = (K - 1)(1/\alpha - 1)$ is selected to satisfy the bound $P(ICD) \leq \alpha$.

The algorithm proceeds as follows:

SE Algorithm

1. Set $n = 0$.
2. Obtain one signal from all area-cells i such that $t_i^{(j)} \neq t_i^{(k)}$ for some configurations $t^{(j)}, t^{(k)}$ in C_n .
3. Compute $\ell_n^{(r)}$ and the ratios $\ell_n^* / \ell_n^{(r)}$, $\forall t^{(r)} \in C_n$.
4. If $\ell_n^* / \ell_n^{(r)} \geq B$, then remove $t^{(r)}$ from C_n .
5. If $|C_n| = 1$, stop and declare the remaining configuration in C_n the *determined configuration*. Otherwise, increase $n \rightarrow n + 1$ and go back to 2.

In step 2 we sample from any area-cell i for which there remain candidate configurations based on different number of TOIs in A_i . Step 3 provides the constraint to the types of problems the algorithm can handle. After the first round of looks, the procedure computes the likelihood for all $\binom{M+A-1}{M}$ configurations. Unless M is small (e.g. less than 5) the number of configurations will make the problem computationally infeasible. For example if the area of interest is a 16×16 grid and there are 10 TOIs, then there are over 10^{16} configurations. In practical situations, however, most configurations may be deemed impossible, so that initially fewer than $\binom{M+A-1}{M}$ configurations exist.

Let the stopping time

$$N_r = \inf \left\{ n \geq 1 : \ell_n^{(1)} / \ell_n^{(r)} \notin (B^{-1}, B) \right\}$$

be the first time the likelihood ratio of Configurations 1 and r exits the interval (B^{-1}, B) . If Configuration 1 is incorrectly eliminated via the SE Algorithm, then there must exist some r such that $\ell_{N_r}^{(1)} / \ell_{N_r}^{(r)} \leq B^{-1}$. The probability that this occurs is $P(\bigcup_{r=2}^K \ell_{N_r}^{(1)} / \ell_{N_r}^{(r)} \leq B^{-1})$. Following Eqs. (4)–(6), with $\alpha = \beta$ and $B = A^{-1} = (K - 1)(1/\alpha - 1)$, we can guarantee the error bound $P(\ell_{N_r}^{(1)} / \ell_{N_r}^{(r)} \leq B^{-1}) \leq \alpha / (K - 1)$. By Bonferroni's inequality, it follows that the incorrect determination probability for the SE Algorithm, $P(ICD)$, will be less than α :

$$\begin{aligned} P(ICD) &\leq P\left(\bigcup_{r=2}^K \ell_{N_r}^{(1)} / \ell_{N_r}^{(r)} \leq B^{-1}\right) \leq \sum_{r=2}^K P\left(\ell_{N_r}^{(1)} / \ell_{N_r}^{(r)} \leq B^{-1}\right) \\ &\leq \sum_{r=2}^K \frac{\alpha}{K - 1} = \alpha. \end{aligned} \quad (7)$$

Hence we conclude that,

Proposition 1. *The SE Algorithm is guaranteed to meet the operator-defined tolerance α for the probability of incorrect determination.*

3.4. Termination time

In this section we show that the SE Algorithm finishes in a finite number of steps with probability 1. To do this we examine the




Fig. 1. Examples of the evolution of the likelihood ratio for the two configuration case.

behavior of the likelihood ratio $\ell_n = \ell_n^{(q)} / \ell_n^{(r)}$ of two arbitrary and different candidate configurations as n increases and determine under what conditions the likelihood ratio exits a continuation region in a finite number of steps.

The continuation region is the interval (A, B) and the stopping time is $N = \inf\{n \geq 1 : \ell_n \notin (A, B)\}$. We will show that if each sensor matrix S_i has no identical rows, then $P[N < \infty] = 1$ for any finite A and B . To do this we examine the log likelihood ratio $\log \ell_n = \log \ell_n^{(q)} - \log \ell_n^{(r)}$. The equivalent continuation region for the log likelihood ratio is (a, b) where $a = \log A$ and $b = \log B$ and the stopping time can be written in terms of $\log \ell_n$ via $N = \inf\{n \geq 1 : \log \ell_n \notin (a, b)\}$.

We will express the log likelihood $\log \ell_n$ as a random walk. That is there are IID random variables Z_j such that

$$\log \ell_n = Z_1 + Z_2 + Z_3 + \dots + Z_n.$$

The values of Z_j can be derived by inspecting Eqs. (2) and (3):

$$Z_j = \sum_{i=1}^A \left(\log p_i(X_{ij}|t_i^{(q)}) - \log p_i(X_{ij}|t_i^{(r)}) \right). \quad (8)$$

The Z_j are independent random variables because each sensor's signals are independent of the signals of other sensors, and a signal from one sensor in one period is independent of a signal from that sensor in another period. The Z_j are identically distributed because the targets are stationary and we assume that the sensor matrices S_i do not change. Therefore, $\log \ell_n = \sum_{j=1}^n Z_j$ is a random walk, and we can appeal to the theory of random walks to prove that $P[N < \infty] = 1$. Because $\log \ell_n$ is a random walk, one of the following four scenarios must occur with probability one (see Theorem 1.2 in Chapter 3.1 of Durrett [2]):

1. $\log \ell_n = 0$ for all n
2. $\log \ell_n \rightarrow \infty$
3. $\log \ell_n \rightarrow -\infty$
4. $-\infty = \liminf \log \ell_n < \limsup \log \ell_n = \infty$

For scenarios 2, 3, and 4, $P[N < \infty] = 1$ because $\log \ell_n$ will almost surely be growing arbitrarily large (or small), and thus $\log \ell_n$ will leave the continuation region (a, b) in a finite number of steps almost surely. To show this we will derive a contradiction. Assume that we have scenario 2, but $P[N < \infty] < 1$. We define the sets D and E :

$$D = \left\{ \omega : \lim_{n \rightarrow \infty} \log \ell_n(\omega) = \infty \right\},$$

$$E = \left\{ \omega : \sup_{1 \leq n} \log \ell_n(\omega) \leq b \text{ and } \inf_{1 \leq n} \log \ell_n(\omega) \geq a \right\}.$$

We have $P[D] = 1$ because we assume scenario 2, and $P[E] = \epsilon > 0$ because we assume $P[N < \infty] < 1$. However, $D \subset E^c$ and therefore $P[D] \leq P[E^c] = 1 - P[E] = 1 - \epsilon < 1$. This is a contradiction and so if we have scenario 2, then $P[N < \infty] = 1$. The proofs for scenario 3 and 4 are nearly identical.

If scenario 1 occurs then $P[N < \infty] = 0$. Since $P[N < \infty] = 1$ if scenario 2, 3, or 4 occurs, if we can prove that scenario 1 cannot happen, then by Theorem 1.2 in Chapter 3.1 of Durrett [2] either scenario 2, 3, or 4 must occur, and therefore $P[N < \infty] = 1$. In scenario 1, $\log \ell_n = 0$ for all n . This will happen if and only if $Z_j = 0$ for $j = 1, 2, \dots$; that is Z_j is identically equal to 0 with probability 1. But this cannot happen if each sensor matrix S_i does not have two identical rows, because then

$$\text{var } Z_j = \sum_{i=1: t_i^{(q)} \neq t_i^{(r)}}^A \text{var} \left(\log p_i(X_{ij}|t_i^{(q)}) - \log p_i(X_{ij}|t_i^{(r)}) \right) > 0, \quad (9)$$

since $\text{var} \left(\log p_i(X_{ij}|t_i^{(q)}) - \log p_i(X_{ij}|t_i^{(r)}) \right) > 0$ whenever $P(p_i(X_{ij}|t_i^{(q)}) \neq p_i(X_{ij}|t_i^{(r)})) > 0$. Eq. (9) implies $P(Z_j = 0) < 1$. Therefore, scenario 1 cannot occur in our model, which means that either scenario 2, 3, or 4 must occur and $P[N < \infty] = 1$.

Since we have shown that the likelihood ratio $\ell_n^{(q)} / \ell_n^{(r)}$ for arbitrary configurations q and r will leave the continuation region in a finite number of steps with probability 1, then this will also hold for the likelihood ratio $\ell_n^* / \ell_n^{(r)}$. This is summarized in the next proposition.

Proposition 2. *If each sensor matrix S_i does not have two identical rows, then the sequential elimination algorithm terminates in finite time with probability 1.*

4. Numerical results

In Section 4.1 we focus on the sequential eliminating algorithm for only one target, and we compare the performance of the sequential eliminating algorithm for one target to other naive algorithms in Section 4.2. We briefly present results for multiple targets in Section 4.3.

4.1. One target scenario

We assume there is a sensor in each area-cell that is available to take one look per time period. To construct the 2×2 sensor matrix S_i for these numerical examples we define $p_i(0|0)$, which corresponds to the upper left hand element of S_i , to be a random uniform value between 0.7 and 1.0, and we assign a random uniform value between 0.8 and 1.0 to $p_i(1|1)$, which corresponds to the lower right hand element of S_i (the exact sensor values used to generate the figures in this section are available from the authors upon request). The quantity $p_i(1|1)$ is the probability the sensor signals that one target is located in the area-cell if there is actually one target in the area-cell. This is known as the *sensitivity* of a sensor. The *specificity* of a sensor corresponds to $p_i(0|0)$, which is the probability the sensor correctly signals that no targets are in the area-cell. The off-diagonal terms of S_i , $p_i(0|1)$ and $p_i(1|0)$, are given once we define the diagonal terms because the rows sum to 1.

Fig. 2 presents the expected number of total looks for the sequential eliminating procedure to terminate as a function of




Fig. 2. Expected number of looks until the sequential eliminating procedure terminates as A increases. There is one target, A ranges from 2 to 256, $\alpha = 0.05$, sensor sensitivity $\sim U[0.8, 1]$, and sensor specificity $\sim U[0.7, 1]$.

the number of area-cells A . We set the error tolerance α to 0.05. The number of area-cells in the AOI varies from $A = 2$ to $A = 256$. We arbitrarily (and without loss of generality) place the target in AC_1 . For all figures in this section we perform 25000 replications for a given set of parameter values to produce an estimated value; in this case that value is the expected number of looks for the procedure to return the determined configuration. The vertical bars represent 95% of the distribution of the number of looks produced during the simulation. Fig. 2 depicts a linear increase in the number of looks as A increases for fixed $\alpha = 0.05$. In this example approximately 3 looks per area-cell are required on average for the algorithm to terminate. This value only increases by about 1 look per area-cell if the total number of looks to terminate is at the 97.5th percentile of its distribution.

In Fig. 3 we examine the effect of varying the error tolerance. We use the same sensor parameters as in Fig. 2, fix $A = 196$, and vary the type-I error probability tolerance α between 0.01 and 0.2. As expected, the number of expected looks decreases with an increase in error tolerance for the sequential procedure. The decrease is more significant for small values of α because the threshold to eliminate configurations varies inversely with α (cf., above Eq. (7)). In the example in Fig. 3, decreasing the error tolerance from 0.2 to 0.01 only requires approximately 1.2 additional looks per area-cell on average.

Finally, the expected number of looks until the sequential eliminating procedure terminates depends upon the sensor parameters in S_i . To determine how improving the sensors' sensitivity and specificity impacts the number of expected looks, we set $A = 196$ and $\alpha = 0.05$ and vary the sensors' parameters. Therefore, the sensor parameters used to create Fig. 4 are different than the values used to produce Figs. 2 and 3. We set the sensor sensitivity to be a uniform random number in $(a, a + 0.05)$ and the specificity to be another uniform random number in $(a, a + 0.05)$ for a ranging between 0.6 and 0.95. As the probabilities that the sensors produce the correct signals increase, the sequential algorithm terminates in fewer steps because there is less uncertainty with the sensors' signals. Increasing the sensors' sensitivity and specificity has a significant impact on the amount of resources (i.e., looks) required to identify the determined configuration. When the sensors produce the correct signal over 95% of the time, then less than 2 looks per area-cell are needed for the procedure to finish. However, when the signals are only slightly better than a coin-flip at 60% the pro-




Fig. 4. Expected number of looks until the sequential eliminating procedure terminates as the sensitivity and specificity increase. There is one target, A is 196, $\alpha = 0.05$, and the sensitivity and specificity range from approximately 0.6 to approximately 0.95.

cedure needs more than an order of magnitude more looks to finish. Fig. 4 illustrates the relationship.

Fig. 4 suggests that one could perform a cost-benefit comparison of several different sensors. “Better” sensors (e.g. those represented on the lower right of Fig. 4) will terminate in fewer looks but they will presumably have a higher purchase cost. There are two primary costs associated with a sensor: the cost to purchase and maintain a sensor and the cost of each look (e.g. the time to process a look). The cost of one look may be quite small, but if that sensor is used in many operations, having a more effective sensor technology may produce significant benefits over time. Another trade-off analysis that could be performed for one particular sensor would be to determine the optimal location on the receiver operating characteristic (ROC) curve. Often the operator has some control over the performance of a sensor; one can increase sensitivity at the expense of specificity (this differs from Fig. 4 where both sensitivity and specificity increase). It would be valuable to know




Fig. 3. Expected number of looks until the sequential eliminating procedure terminates as α increases. There is one target, A is 196, sensor sensitivity $\sim U[0.8, 1]$, sensor specificity $\sim U[0.7, 1]$, and α ranges from 0.01 to 0.2.




Fig. 5. Expected number of looks until the sequential eliminating procedure terminates as the sensitivity and specificity vary. There is one target, A is 196, $\alpha = 0.05$, and each sensor has the same sensitivity and specificity. The darker shades represent a larger number of looks.

which combination of sensitivity and specificity would produce the smallest expected number of looks. A full cost-benefit analysis of sensor technology is outside the scope of this paper (and much of the relevant information would be classified or difficult to obtain). However we provide one figure that would be an essential component of that analysis once one has the relevant cost and sensor data. Fig. 5 presents a heat map that plots the expected number of looks as a function of the sensitivity and specificity (we assume the sensor characteristics are identical across area cells). This figure illustrates that there are primarily eight regions in the sensitivity/specificity space where the expected number of looks is roughly constant (technically there is a very small ninth region in the upper right hand corner for near perfect sensors that is difficult to see on the figure). The expected number of looks in the first region in the lower left-hand corner of Fig. 5 is approximately 6500 and that value drops by at least 25% per region as the sensor parameters move to the northeast until the value is less than 500 in the eighth region. Operationally, Fig. 5 suggests that, in terms of expected number of looks until termination, the marginal values of the sensitivity and specificity are similar when their values are alike. The more dissimilar their values, the greater the marginal value of the largest parameter becomes. The structure of Fig. 5 would be useful in the evaluation of sensor technologies. For most situations a small improvement in sensor performance will not justify the additional cost. However, if current sensor technology is on the boundary of two regions, the extra investment will significantly reduce the number of looks.

4.2. Comparison with other target-location algorithms

In this section we compare the performance of the sequential eliminating procedure with two other algorithms to locate targets in an AOI; the measure of performance is the observed error rates. Recall that the goal of the sequential eliminating procedure is to indicate target location with an accuracy rate guaranteed to meet operator-specified error tolerances, within a reasonable number of expected looks. One possible measure of performance is the amount of slack between the error tolerance and the observed error rate. Intuitively, the less slack, the fewer number of expected looks would be required. However, if a method with the same expected number of looks as the sequential eliminating procedure exhibits a larger observed error rate, it is reasonable to state that the sequential eliminating procedure is more efficient than such a method. Alternatively, one could invoke a method exhibiting the same achieved error rate and compare the expected number of looks, but we choose the former scheme for ease of computation and illustration.

In both of the alternate methods we consider, we allocate the same expected number of total looks as the sequential eliminating procedure. For given values of S_i and α , let $\bar{\tau}_A$ be the total average number of looks for the sequential eliminating procedure to complete with an AOI consisting of A area-cells. We compute the average number of looks numerically in Figs. 2–4. In the naive procedure, we allocate the same number of looks to each sensor so that each area-cell receives $\bar{\tau}_A/A$ looks. In a slightly more sophisticated procedure, we allocate the looks to area-cells based on how effective their sensors are. Looks are allocated in proportion to the inverse of (sensitivity - (1-specificity)). If this quantity is large (that is the sensor has high sensitivity and specificity) the sensor is efficient and needs less looks to provide useful information to the operator. If the difference between sensitivity and the Type-II error of the sensor is small then it is difficult to interpret one signal with great certainty, and therefore more signals from those sensors are collected. If AC_i should receive a fraction f_i looks according to this procedure, then $f_i \bar{\tau}_A$ looks will be allocated to that area-cell.

In both the procedure that equally allocates looks to every area-cell and the procedure that allocates looks in proportion to the inverse of sensor effectiveness, the calculated looks-per-cell will likely not be integral. We round the calculated value *up* to the nearest integer, and therefore these two alternate procedure will have more total looks than the sequential eliminating procedure. After the deterministic allocation of looks to the area-cells, the configuration with the largest likelihood is declared the determined configuration for each of the algorithms. Unlike the sequential eliminating procedure, neither configurations nor area-cells are eliminated during these two alternate procedures.

In Fig. 6 we compare the observed error rates of the sequential eliminating procedure with that of the alternate models. We set $\alpha = 0.05$ and vary the number of area-cells from $A = 2$ to $A = 256$. As before, the target is located in AC_1 and the sensor parameters are the same ones used produce Fig. 2. Furthermore, for the sequential eliminating procedure we use the same data that produced Fig. 2. We do not include vertical bars in Figs. 6–8 because we are only computing one realization of the error rate (as opposed to 25000 realizations of the number of looks in Figs. 2–4). As constructed, the observed error rate of the sequential eliminating procedure is less than the error bound α (significantly so in this case, due to the slack introduced by Bonferroni's inequality, cf. Eq. (7)). For the same average number of looks at each value of A , the observed error rates for the two alternate methods are considerably larger for all values of A (indeed, the alternate error rates exceed the tolerance α for most values of A). In many situations the algorithm that allocates looks inversely to sensor effectiveness produces more errors than the algorithm that allocates an equal number of looks to each area-cell. By potentially allocating too many looks to too few area-cells the former algorithm does not collect enough information from the remaining area-cells to accurately produce the correct configuration.

We next examine the effect of varying α with $A = 196$. Fig. 7 shows that an increase in α appears to cause a linear increase in the observed miss rate. The observed miss rate for $\alpha = 0.2$ is still much less than 0.05. The bound derived in (7) is not tight, and as Fig. 7 illustrates there is a large amount of slack between the desired error rate and the observed error rate. This suggests there is potential to improve the sequential eliminating procedure by terminating after a fewer number of looks. The difference between the observed error rate for the sequential eliminating procedure




Fig. 6. Observed error rates of the target-location algorithms as A increases. There is one target, A ranges from 2 to 256, $\alpha = 0.05$, sensor sensitivity $\sim U[0.8, 1]$, and sensor specificity $\sim U[0.7, 1]$.




Fig. 7. Observed error rates of the target-location algorithms as α increases. There is one target, A is 196, sensor sensitivity $\sim U[0.8, 1]$, sensor specificity $\sim U[0.7, 1]$, and α ranges from 0.01 to 0.2.




Fig. 8. Observed error rates of the target-location algorithms as the sensitivity and specificity increase. There is one target, A is 196, $\alpha = 0.05$, and the sensitivity and specificity range from approximately 0.6 to approximately 0.95

and the two alternate procedures is significant: the error rates for the latter two start at over 15% and increase to over 25%.

Finally we present the companion to Fig. 4 in Fig. 8 for $\alpha = 0.05$ with $A = 196$. As the sensors improve and they more accurately signal the true nature of the area-cells, the error rates of all the procedures generally decrease. Although the observed error rates of the two alternate procedures are still greater than $\alpha = 0.05$, except when the sensitivity and specificity are greater than 0.95. It is not surprising that all three algorithms have relatively low error rates when the sensitivity and specificity are both large because often one or two looks per area-cell are sufficient to produce a determined configuration. However, when sensors are not effective and the difference between the probability of a true positive and a false positive gets closer to zero, the dominance of the sequential eliminating procedure over the alternate methods is staggering. The alternate methods produce nearly as many incorrect configurations as true configurations, whereas the sequential eliminating procedure is still comfortably below the $\alpha = 0.05$ error rate. Often,

depending upon the circumstances (e.g. the sensor technology, environment, nature of the targets, and ability of the operators), the sensors' signals will be unreliable and the sensitivity and specificity will be much less than 0.9. These are the situations where the sequential eliminating procedure can produce enormous efficiency gains.

4.3. Multiple targets

The results for multiple targets are similar to the case of one target. Fig. 9 presents how the expected number of looks for the sequential eliminating procedure to terminate varies as the number of targets, M , increases. We set $\alpha = 0.05$, $A = 10$, and vary M from 1 to 10. We define the GTC to be one target in the first M area-cells. There are many ways the sensors' capabilities can change as the number of potential targets increases. We examine two cases. In one case a sensor's ability to signal the correct number of targets in the area-cell is a uniform random variable between 0.5 and 1 (these are the diagonal values in S_i). This means that, regardless of the number of targets, the sensors produce a correct signal at least 50% of time. The remaining sensor parameters are uniformly generated to ensure that the rows are pmfs (i.e., the rows sum to 1). We also examine a case where the probability a sensor produces a correct signal decreases as the number of targets increases. For simplicity, in this case we assume that the off-diagonal elements of S_i are identical and that the diagonal elements are a multiple of the off-diagonal terms. We assume the diagonal element of each row of S_i is a factor of 4 larger than an off-diagonal element because this corresponds to a sensitivity and specificity of 0.8 when there is one target. We further assume the parameters of every sensor are the same, and therefore the off-diagonal elements of S_i equal $1/(M + 4)$ and the diagonal elements of S_i equal $4/(M + 4)$.

Adding additional targets rapidly increases the number of feasible configurations. With one target there are A configurations (i.e., the target can be in one of the area-cells). However, with 5 targets in 10 area-cells there are over 2000 configurations. Not only are there more configurations to consider, but the threshold to eliminate configurations increases with the number of feasible configurations (cf., above Eq. (7)). It might be expected that, as with the number of configurations, the number of looks would increase exponentially as we increase M . However, for the two cases we




Fig. 9. Expected number of looks until the sequential eliminating procedure terminates as M increases for two different categories of the sensor matrices S_i , $A = 10$, $\alpha = 0.05$ and M ranges from 1 to 10.

examine in Fig. 9 this does not occur. In fact when the sensor correctly identifies the target over 50% of the time, the average number of looks is relatively constant as M increases. This is because the probability of returning a specific incorrect signal is much less than when there is only one target. Therefore, the configurations that do not match most of the sensor signals will quickly be eliminated. When the probability that a sensor produces the correct signal decreases as M increases, the number of looks appears to increase linearly. In this case, additional targets cause the sensors to produce more incorrect signals, and therefore to overcome this uncertainty the sensors must take additional looks. In the example in Fig. 9, every additional target requires approximately 4 more looks per area-cell to find the determined configuration.

When multiple targets are located in the AOI, the observed error rates are similar to the results presented in Figs. 6–8. The sequential eliminating procedure produces errors at a rate much less than α and it dominates the two alternate procedures. However, the error rates of the two alternate procedures are smaller than when there is only one target. Because of the large number of configurations, the threshold to eliminate configurations is greater than it needs to be to achieve the desired error rate, which causes considerable slack between the observed and desired error rates.

5. Concluding remarks

The sequential eliminating procedure collects signals from sensors and determines how targets are located in area-cells in an area of interest. This procedure provides efficient results guaranteed to meet desired error rate bounds. The key assumptions are that targets are nonreactive and static, or that the targets remain in the same area-cell during the time-scale of interest.

We show that two naive, non-eliminating methods demonstrate consistently higher error rates for the same total number of looks for all appropriate values of number of area-cells, error tolerances, and sensor parameters. There is also no guarantee that the naive procedures' results meet the operator-prescribed error tolerance. In fact, for most of the examples we analyze, their error rates are higher than the tolerance. It is therefore reasonable to conclude that the sequential eliminating procedure is more efficient than the naive allocation methods. The sequential eliminating procedure is especially superior to the other two methods when the sensors have lower sensitivity and specificity (cf., Fig. 8). The number of looks scales linearly with the the number of area-cells and targets for the examples in Section 4. This suggests the sequential eliminating procedure can be utilized in scenarios involving a large region with many targets.

The sequential eliminating procedure can be useful even if it cannot run until completion. This may occur if there are a limited number of looks available due to resource or time constraints. The sequential eliminating procedure can run until the look limit is reached, and then with probability $1-\alpha$ the targets are located according to one of the remaining candidate configurations. Further investigation by different search assets (e.g. interceptor

teams) may be required to determine the location of the targets. The sequential eliminating procedure reduces the possibilities to the remaining candidate configurations. If one determined configuration must be returned at the end of the procedure, then the configuration with the largest likelihood can be chosen.

Finally, the analysis in this paper presents the worst case performance in some sense. The sequential eliminating procedure terminates only when one configuration remains, and we consider an incorrect determination to occur if the determined configuration is not exactly the ground truth configuration. The procedure could return a set of candidate configurations rather than one determined configuration, and for many applications several acceptable configurations may exist. For example the goal may be to determine which area-cells contain any targets; the exact distribution of the targets in the area-cells may not be important. Our analysis provides an upper bound on the error rate and number of looks if there are multiple acceptable configurations.

References

- [1] S.J. Benkoski, M.G. Monticino, J. R. Weisinger, A survey of the search theory literature, *Naval Research Logistics* 38 (4) (1991) 469–494.
- [2] R. Durrett, *Probability: Theory and Examples*, Brooks/Cole – Thomson Learning, Third ed., Belmont, CA, 2005.
- [3] W. Feller, *An Introduction to Probability Theory and Its Applications*, Third ed., vol. 1, John Wiley & Sons, New York, 1968.
- [4] J.B. Kadane, Optimal whereabouts search, *Operations Research* 19 (1971) 894–904.
- [5] B.O. Koopman, *Search and Screening*, Center for Naval Analyses, Alexandria, VA, 1946.
- [6] M. Kress, J.O. Royset, Aerial search optimization model (ASOM) for UAVs in special operations, *Military Operations Research* 13 (2008) 23–33.
- [7] M. Kress, R. Szechtman, J.S. Jones, Efficient employment of non-reactive sensors, *Military Operations Research* 13 (2008) 45–57.
- [8] E.L. Lehmann, *Elements of Large-Sample Theory*, Springer-Verlag, New York, 1999.
- [9] G.J. Malone, *Ranking and Selection Procedures for Bernoulli and Multinomial*, Ph.D. Thesis, Georgia Institute of Technology, 2004.
- [10] Y. Mei, *Asymptotically Optimal Methods for Sequential Change-Point Detection*, Ph.D. Thesis, California Institute of Technology, 2003.
- [11] P. Owen, R. Martin, T. Carriger, Shadowing IEDs, *Unmanned Systems* 2 (2005) 13–16.
- [12] D.N. Reber, *Optimized Routing of Unmanned Aerial Vehicles for the Interdiction of Improvised Explosive Devices*, Master's thesis, Naval Postgraduate School, 2007.
- [13] H. Sato, J.O. Royset, Path optimization for the resource-constrained searcher, *Naval Research Logistics* 57 (5) (2010) 422–440.
- [14] D. Siegmund, *Sequential Analysis: Tests and Confidence Intervals*, Springer-Verlag, New York, NY, 1985.
- [15] N-O. Song, D. Teneketzis, Discrete search with multiple sensors, *Mathematical Methods of Operations Research* 60 (1) (2004) 1–13.
- [16] T.H. Chung, M. Kress, J.O. Royset, Probabilistic Search Optimization and Mission Assignment for Heterogeneous Autonomous Agents, in: *Int'l. Conference on Robotics and Automation*, 2009, pp. 939–945.
- [17] K.P. Tognetti, An optimal strategy for whereabouts search, *Operations Research* 16 (1968) 209–211.
- [18] K.E. Trummel, J.R. Weisinger, The complexity of the optimal searcher path problem, *Operations Research* 34 (2) (1986) 324–327.
- [19] A.R. Washburn, *Search and Detection*, fourth ed., Institute for Operations Research and the Management Sciences, 2002.
- [20] J.R. Wieland, B.L. Nelson, An odds-ratio indifference-zone selection procedure for Bernoulli populations, in: *Proceedings of the 36th conference on Winter simulation*, Winter Simulation Conference, 2004, pp. 630–638.